Chapter 6

AUCTIONS AND PRICING IN E-MARKETPLACES

Wedad Elmaghraby
Georgia Institute of Technology
Atlanta, GA 30332
wedad@isye.gatech.edu

1. B2B E-Marketplaces

By offering high-speed communication and tight connectivity, advances in information technology have opened new venues for companies to create flexible supply chains. Today, many companies, from the electronics, pharmaceutical, to the automotive industry, are focusing on their core competencies and outsourcing significant portions of their business operations. As supply chains become more decentralized upstream, the pricing of intermediate goods is no longer a formality used between departments - but rather the key to a company's survival. Strangely enough, while companies have spent millions of dollars to help them reduce their operating costs, the majority continue to use crude techniques in deciding what price to charge for their products [Anthes, 1999].

The increased adoption of revenue management pricing policies, along with the emergence of Internet enabled marketplaces, referred to as e-marketplaces, is one of two separate but interrelated phenomenom that is helping to change business practices. Sophisticated capacity allocation strategies, the essence of current revenue management tools, are burgeoning on the business-to-consumer (B2C) front. Designed for markets with perishable products and limited supply, revenue management tools aid companies in *dynamically* selecting the optimal mix of products to offer to customers. This is done by allowing customers to make advance bookings for different products (fare classes) and dynamically changing the availability of each product [McGill and van Ryzin, 1999] [Phillips, 2003]. While revenue management tools focus on capacity allocation decision, they have generally ignored related pricing decision. That is, revenue management tools take the prices for different products as given,[1] rather than simultaneously determining the optimal price for each product with the capacity allocation.

E-marketplaces allow companies to go one step beyond traditional revenue management techniques, by offering a natural medium to optimally develop sophisticated pricing policies.Via e-marketplaces, suppliers (and buyers) can reach larger markets, dynamically change prices as the need and opportunity arise, and, most importantly, gain vital information about demand elasticity to incorporate into their pricing decisions.

While companies have embraced revenue management and e-marketplaces as essential to their future growth on the B2C front, they have been slow to alter their current business-to-business (B2B) practices. In particular, companies have been slow to use e-marketplaces as a transaction medium, due to the slow pace of technology diffusion and the complexity of most B2B transactions.[2] While B2C transactions can generally be characterized as involving spot transactions of a known (small) size, B2B transactions are generally governed by contracts over an extended period time with uncertain demand and entail identifying multi-attributes in addition to price, e.g., transportation, quality, ordering and replenishment procedures, etc. Therefore, except for the sale of some commodity parts and excess inventory, e-marketplaces have played a minimal role in most B2B transactions [Latham, 2000]. The fact that many companies do not currently use the Internet to help them with their B2B pricing decisions does not mean that there is no future potential nor enormous gains from doing so. It is anticipated that the value of products and services sold in B2B markets through electronic channels in the United States will reach $4.7T by 2005 [AMR, 2001].

The purpose of this paper is to bring to the surface the current pricing practices in B2B e-marketplaces and the challenges in successfully implementing pricing policies in practice. In particular, we focus on two pricing strategies that are natural candidates for B2B e-marketplaces; *auctions* and *precision pricing* (referred to as 3^{rd} degree price discrimination in the economics literature). Auctions, whereby prices are determined via a bidding process, have already made some headway in B2B e-marketplaces. For example, retailers in footware, home products and fashion are using GlobalNetXchange private auction exchange [Rosenblum, 2002], auto manufacturers are using Covisint's auctions capabilities [Prouty, 2001] and GE uses its own Global Exchange Services [Barling, 2001] to help procure goods more effectively from suppliers. Precision pricing entails quoting each customer-type a different take-it-or-leave-it price based on relevant and observable characteristics. Precision pricing's presence is more sparse in B2B e-marketplaces, but is increasingly being adopted in off-line sales channels with an eye towards adoption in e-marketplaces.

By presenting case studies of companies that design and implement auctions and precision pricing tools, we hope to gain a better understanding as to the future role and contributions the ORMS community can have in developing and executing these pricing policies. We begin by providing an overview of the pricing policies used currently in e-marketplaces.

2. Current State of Pricing in B2B Marketplaces

To better understand the current pricing practices and future potential for pricing policies in e-marketplaces, it is helpful to categorize market environments according to (i) the *size* and *frequency* of the transaction, (ii) the ability for arbitrage and (iii) the information asymmetry in the marketplace. The *size* and *frequency* of the transaction determines whether or not it is cost-effective to tailor a pricing policy to that specific transaction. For large and infrequent transactions, it is typically worthwhile to take the time and money to gather information about the marketplace and design a pricing policy accordingly. On the other hand, for small and frequent transactions, it is too costly and inefficient to specifically design a pricing policy to the particular transaction. Rather, a pricing policy should dynamically response to underlying changes in the marketplace in a (semi-)automated manner. The ability for *arbitrage* and the *information asymmetry* in the marketplace critically influence the degree to which the seller can practice price discrimination, i.e., charge different prices to different customers. It is difficult for the seller to price discriminate if a product can be easily sold in a resale market (for the market segment who buys at a lower price can sell to other market segments), if customers can easily acquire information concerning the true value/cost of a product, or if the seller has little information about customer valuations. However, if different market segments can be identified across which resale is difficult, then a seller can charge different prices to each group.

During the late 1990's, a host of B2B e-marketplaces emerged in both the service and goods industries. Companies in the transportation, plastics, chemical, and paper were among the first to build and use B2B e-marketplaces; in these industries we saw the creation of Cargonow.com, The PlasticsExchange.com, ChemConnect.com and Paperloop.com.

One of the first uses for these e-marketplaces was for the sale of excess inventory (region A in figure 6.1). Sellers in these markets who found themselves with unsold units eagerly turned to e-marketplaces as a medium via which they could reach a larger and previously untapped market and derive a higher salvage value for their items. In effect, e-marketplaces served as ‘garage sales’ for these sellers. These garage sales were initially operated as a bulletin board, whereby sellers post their products and request price. With a bulletin board, a transaction takes place as soon as one customer who is willing to pay the seller's requested price contacts the seller. This method, while reaching a larger customer base, does not allow sellers to use competition to drive up the selling price of their products. Therefore, the second type of pricing policy used in these garage sales is an auction. The most commonly used auction format is an ascending price English auction, whereby the auction terminates when bidding ceases and the product(s) is awarded to the highest bidder(s) at the highest bid price. In an ascending price English auction, interested buyers submit increasingly higher bids. A key attribute of these garage sales (posted price or auction) is that they are managed by the seller, i.e., the e-marketplace merely serves as a host site.

Alongside the sale of excess inventory via garage sales, e-marketplaces began to build and run double auctions for commodities, similar in nature to the Chicago Mercantile Exchange (region B in figure 6.1). It is often difficult to limit arbitrage with pure commodities, therefore a pricing mechanism that yields the same price to all customers at any one point in time is most appropriate. In a double auction, buyers (sellers) submit their bid (ask) prices indicating the price they are willing to pay (receive) for a specified quantity. The market maker aggregates the supply and demand and determines a market price that clears the market and reflects the 'true value' of the product. With commodity products, there are few if any buyer-supplier specific attributes aside from price that influence the transaction decision. This is in sharp contrast to customized products, for which identifying compatible suppliers as well as supplier specific characteristics such as quality and service can play an equally if not more important role than price. Double auctions are actively managed by the e-marketplace.[3]

The third type of pricing policy now entering e-marketplaces are customized auctions (region C in figure 6.1). An estimated 80-90% of all B2B transactions are done via long-term contracts [Economist, 2001]. Given the vast sums of money that are contracted upon and the multi-attribute nature of many B2B transactions, it is clear that these less frequent and larger transactions cannot be handled by standard double auctions nor self-managed garage sales. Rather, they need a pricing mechanism that is rich in its dimensions so that it can be specifically tailored to the situation at hand. That is, the pricing mechanism should be more 'hands on' and allow for the proper specification of the product, be flexible enough to reflect various market participant characteristics and adapt to changing circumstances in the marketplace. Customized auctions fit the above description and are being used to help buyers establish relationships with suppliers. This vast sum of untapped money has not gone unnoticed by makers of e-marketplaces; customized procurement auctions for large and infrequent transactions require significant auction and industry expertise in their design. In response to this need, we are beginning to see companies such as FreeMarkets, Logistics.com[4], ChemConnect.com and CombineNet offer companies support in the design and implementation of these auctions.[5]

Figure 6.1 illustrates the placement of each of the current pricing practices along these three market dimensions; arbitrage, information asymmetry and frequency of transactions, while figure 6.2 provides an example of the pricing policies for the transportation industry. In the

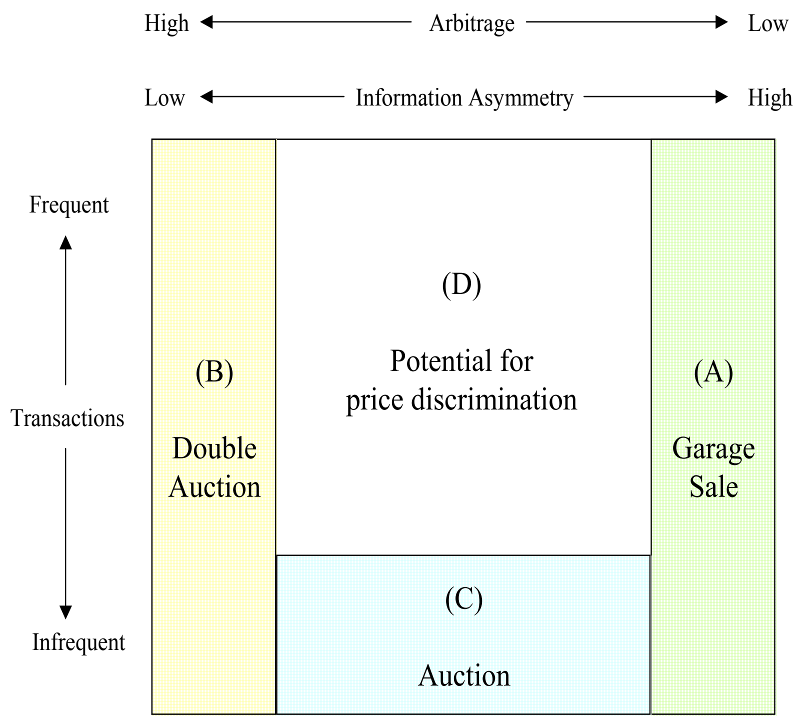

Figure 6.1. Current pricing practices in B2B e-marketplaces.

transportation industry, Logistics.com aids shippers who wish to procure long-term transportation contracts with multiple carriers via combinatorial auctions.[6] With their OptiBid procurement system, Logistics.com aids shippers design and run these large (fully personalized) procurement auctions, by helping them create appropriate transportation networks, send auction information to and receive bids from the carriers, and optimize which carriers to select given the submitted bids (see region A of figure 6.2). Although most shippers and carriers enter into long-term contracts for the bulk of their business, imbalances in supply and demand often occur, giving rise to either empty carrier capacity or excess shipments that need to be transported. The National Transportation Ex-

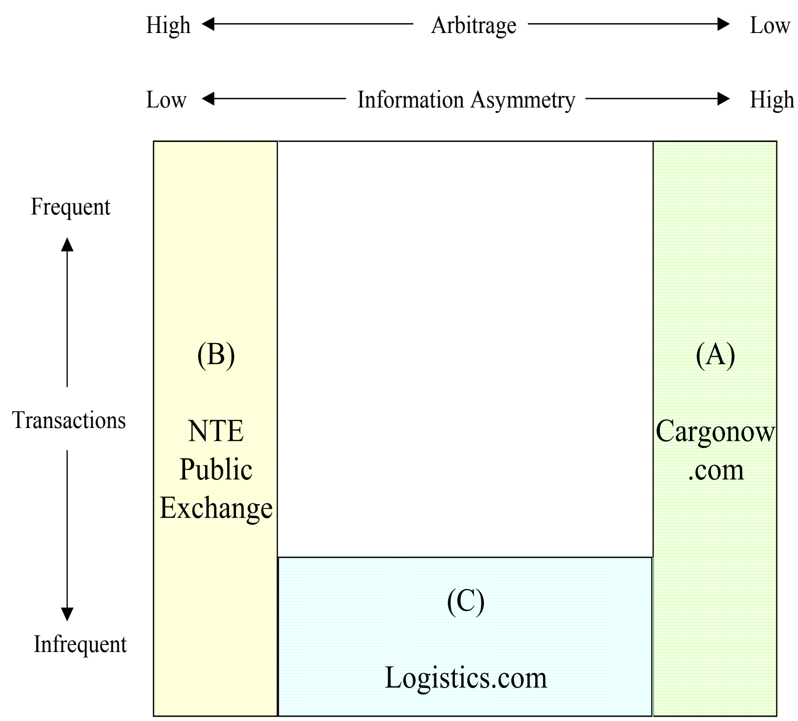

Figure 6.2. E-Marketplaces currently being used in the transportation industry.

change (NTE) helps to fill in these imbalances in supply and demand by offering a dynamic, automated transportation marketplace, NTE Public Exchange, where carriers can submit empty or partially filled capacity, and shippers can submit excess demand (region B in figure 6.2). The NTE Public Exchange is in essence a double auction over multiple products (a product being an origin-to-destination route); it matches the submitted demand and supply and determines a transaction price for each match obtained. A shipper or carrier that prefers to post its own spot requirements (and circumvent the transaction fee associated with using a public exchange such as NTE's) can use Cargonow.com's website (region A in figure 6.2).

Currently, e-marketplaces have captured the extremes of the market: There are auctions (region C) for large and infrequent transactions, double auctions (region B) for the sale of high arbitrage/ low information asymmetry (commodity) products and garage sales (region A) for the sale of high information asymmetry products such as excess inventory. What we do not observe are e-marketplaces and pricing policies for medium sized transactions where arbitrage is difficult and the seller has some information about customer valuations (region D).

In selecting a pricing policy for region D, a seller must tradeoff two main factors: (i) the *simplicity* to develop and administer the pricing policy and (ii) the *profitability* of the pricing policy. As a general rule, the more complex is a market environment, the greater the potential rewards from the use of a sophisticated pricing strategy. While double auctions for commodities and 'garage sales' are relatively simple to develop and administer, they do not take advantage of customer heterogeneity and seller information by price discrimination. One particular pricing policy that lends itself to region D market transactions is *precision pricing.* Precision pricing allows a seller to tailor his prices to different segments of his market, thereby exploiting differences in customer valuations by charging them different prices.

Where a company is able to use individualized prices and/or is able to effectively segment customers, we foresee large infrequent transactions remaining in the realm of auctions, while precision pricing will be employed for smaller more frequent transactions. Precision pricing, along with customized auctions, offer its user the promise of increased profitability. But the promise of increased profitability does not come without serious challenges. A key lesson learned from early entrants in B2B e-marketplaces is that the interconnectiveness of the Internet is not sufficient to make on-line pricing a successful venture. What is necessary is the know-how behind selecting the appropriate pricing strategy and the expertise to properly design and implement it.

To better understand the design dimensions and implementation challenges present in the use of auctions and precision pricing, we present below a few pricing 'success stories'. In section 3 we discuss the multiple decisions that FreeMarkets, a premiere customized auctions company, faces when designing a procurement auction.[7] In section 4, we discuss the design and use of a bidding support software, developed by Manugistics, to aid bidders participating in a request for quotes. In section 6, we present a description of Manugistic's precision pricing tool, the challenges in properly implementing third degree price discrimination. We conclude each topic with future directions of research for the field of ORMS.

3. Customizing Auctions - A Case Study of FreeMarkets

In the 'new' economy, customized auction market makers see themselves as transformers of supply chains. Where previously supply chains were rigid with fixed suppliers and uncompetitive procurement practices, auction market makers infuse change by breaking down barriers that obscure information, bringing new suppliers, technology and expertise to monitor the supply chain, and streamlining the procurement process, thereby allowing it to be more agile and transparent, and further rewarding efficient and competitive suppliers. One of the leading customized auction market makers in B2B e-marketplaces is FreeMarkets.

As noted by Jason Busch, Senior Manager of Strategic Marketing at FreeMarkets,[8]

"The Center for Advanced Purchasing Studies and AT Kearney did a study of all the potential ways that savings can be generated from procurement. Small gains can be achieved by running operations more efficiently. There are savings to be had in reducing inventory, cutting purchasing headcount, and ensuring all MRO buys are made on contract. But the real bang for the buck is reducing the costs of direct and indirect purchases in an organization...73% of purchasing costs are in sourcing; i.e., in what is being bought...Every $1 saved in sourcing is worth an estimated $5-$25 in increased sales."

Started in 1995, FreeMarkets was one of the first companies to offer customized auction *services* to industrial customers on a large scale. The emphasis on 'services' is clear to anyone who has designed or participated in a large industrial auction. The actual auction is the final step of the auction design process; in order for it to be a success, a substantial amount of information and work must be done beforehand in preparing both the buyers and suppliers. FreeMarkets' main focus is on helping companies through the entire sourcing process, starting from adequately specifying what product(s) the buyer wishes to purchase, identifying potential suppliers on a global scale, and designing the appropriate auction mechanism for the market transaction. Their primary product, Full Source, is a full-service product which includes FreeMarkets expertise in 8 vertical industries (Electronics, Engineering and Construction, Metals, Paper & Packaging, Plastics, Raw Materials, Maintenance Repairs Operations & Services, and Transportation), as well as an auction software which brings together supplier information and market making services in e-space. Currently, FreeMarkets has 125 customers, including Tier 1 and 2 automotive companies and oil companies such as BP Amoco.

While there are several sizeable challenges to making one of their auctions work, FreeMarkets has enjoyed considerable success in the auctions business. To their customer's procurement process they bring transparency, a list of 14,000 suppliers all over the globe, and cost competitiveness. In addition, they bring a reduction in the time it takes to procure. In a world where time is money, FreeMarkets' ability to significantly reduce procurement cycle times allows their customers to bring products to their market more quickly. For example, one of FreeMarkets' clients, Emerson, was able to reduce its procurement cycle time reduction from 12-21 weeks down to 8-12 weeks, with the use of FreeMarkets' Full-Source. The savings from these cycle time reductions can range anywhere from 5-15% based on the industry and products being auctioned.[9]

Before holding an on-line auction event, FreeMarkets must go through four critical preparation steps: (*i*) fully specifying all dimensions to a buyer's entire order, (*ii*) standardizing information across suppliers, (*iii*) monitoring and identifying supply market trends, and (*iv*) selecting an appropriate auction format. Surprisingly, many buyers will initiate a procurement process without a clear and definite idea of what it is they wish to purchase. While they may know the general characterization of the product, they often have not thought through clearly all the products' dimensions and supply specifications, such as delivery dates, order size, and required service levels. An ill-defined order can lead to misunderstanding and uncompetitive bidding. In order to make their sourcing process a success, FreeMarkets spends a considerable amount of time *detailing* the design of the product, i.e., pinpointing the exact specifications of the product the customer wishes to procure. On the opposite side of the market, the supplier base is often quite diverse, be it with respect to geographic location, technology, production capabilities and/or quality.[10] Therefore, FreeMarkets must *standardize* the product information, e.g. by defining a minimum threshold level of service and scope so as to level the playing field for the suppliers and allow the bids in the auction to be meaningful. Given the dynamic nature of the industries in which FreeMarkets operates, they carefully *monitor* and *identify* market trends so as to identify eligible suppliers to invite to submit bids and develop a better handle on their cost structures and capacity availability. For example, suppliers may change in their product/service offerings, the geographical markets in which they are operating, and their performance levels. Finally, once the product and potential suppliers have been specified, auction makers, who understand the supplier market, work with engineering experts to select the optimal auction format. With various options with respect to auction format (e.g., open or sealed-bid, qualifying round), bid format, and feedback format (discussed further below), FreeMarkets has over 30 auction formats. Their strength and success lies in their ability to properly match the market setting with an auction format.

3.1 Auction Formats

If an auction is designed "properly", it holds the promise of achieving a buyer's objectives of cutting down procurement costs and/or awarding business to "healthy" suppliers with whom she can enter into long-term strategic partnerships. Many of the companies (including United Technologies, Quaker Oats, General Electric, and Owens Corning) who have turned to FreeMarkets to help them in reducing their procurement costs have reported significant savings as a result of the on-line customized auction. In order to attain these savings, FreeMarkets has had to tailor the format of each auction to the situation at hand. There are four main auction design dimensions that must be carefully selected for the success of an auction: (*i*) how to *bundle* the buyer's demand and construct bidding lots, (*ii*) how to *sequence* the lots' auctions, (*iii*) what type of feedback to provide to bidders during the auction, referred to as *feedback to the marketplace* and (*iv*) what type of bid format, or *bid calculation* form, to use.

Given the multi-unit/product nature of most of their buyer's demand, FreeMarkets must decide which units/products to *bundle* together and have suppliers bid on as an entire package, or lot. It is often advantageous to bundle pieces of an order together either because the items experience synergies in production or because packaging allows a desirable job to be bundled with an undesirable one.[11] Due to supplier capacity constraints FreeMarkets can not typically bundle all of a buyer's order together into one package. FreeMarkets must trade off increasing the size of the bundle and allowing for more economies of scale in production with reducing the number of suppliers who have sufficient capacity to supply the entire bundle and hence can compete in the auction. Asymmetric supplier capacities further exacerbate the bundling decision. In order to determine the best number and size of lots to create, the auction maker must have a solid understanding of the bidding suppliers' production capabilities (technology), available capacity, and cost structure.

Once the components of the bundles are determined, FreeMarkets must decide on the *sequencing* of the auction. The lots on which suppliers are submitting bids typically consist of large complex orders that require the supplier to collect a substantial amount of information and resources so as to submit an appropriate bid. FreeMarkets has observed that, typically, suppliers dislike participating in more than one auction simultaneously; they prefer instead to focus their bidding on one bundle at a time, with the winner in each auction being notified at the end of the auction. As a result, FreeMarkets staggers the closing of auctions and faces the challenge of deciding in which order to sequence them. Although they do not have a general policy for ordering, FreeMarkets typically closes the largest (major) lots last.[12]

The amount of information that is observed by the auction participants can play a pivotal role in combating supplier collusion during an auction and getting reluctant suppliers to participate in the on-line auction. Therefore, the *feedback to the marketplace* is a critical auction dimension. The majority of FreeMarkets' on-line auctions are open or iterative auctions as opposed to sealed-bid auctions. The preference for open/iterative auctions stems from the belief that the psychology of an open auction induces suppliers to bid more aggressively. In special cases where FreeMarkets is procuring from a very small and collusive supplier base, e.g. the sugar market where there are only 2 main suppliers, FreeMarkets reverts to the use of a sealed-bid auction in order to reduce the information exchange and the ability of suppliers to collude and maintain high prices in the auction.

In open/iterative auctions, FreeMarkets has four feedback formats; (*i*) full disclosure, (*ii*) rank, (*iii*) next horse, and (*iv*) starting gate. Ideally, FreeMarkets would like to always provide the bidders with *full disclosure*, where bidders can see all the current bids as well as who has placed them.[13] Their preference for full disclosure is based on their experience that, in a competitive supplier market, an auction performs best when as much information as possible is made known to the bidders. Unfortunately, many supplier markets, such as metals, ocean freight, construction and air parcel, are not very competitive. In these markets, suppliers are aware of the opportunities for collusion in a full disclosure auction and the possible repercussions associated with participating in an auction.[14] [15] Therefore, in an effort to combat collusion and encourage supplier participation in their auctions, FreeMarkets uses an alternative feedback format, a *rank* feedback auction. A rank auction allows all the bidders to see their and their competitors' rankings, but not the submitted bids. Rank and next horse feedback are used when the buyer is purchasing from a supplier market that is particularly disparate in cost or quality. Under *next horse feedback* format, bidders can only see the bid that is directly in front of (below) them. In addition, FreeMarkets will often usually use a *starting gate feedback* form to supplement the above three feedback forms. With starting gate feedback, bidders are disallowed from seeing anything until they place a bid. This is used to avoid suppliers from promising that they will participate in the auction, only to use the auction as a means to gain information about the marketplace without ever placing a bid.

The selection of *bid calculation* format depends on whether the suppliers are competing to sell a commodity part for which the identity of the supplier is of little importance, or if the suppliers, and their products, are heterogenous. The four bid calculation formats that are used most frequently are (*i*) index, (*ii*) descending, (*iii*) transformational and (*iv*) net present value (NPV). An *index* bid is commonly used in the sourcing of a raw material. Often raw materials are traded on a public exchange, such as NYMEX, and hence have a publicly observable market price. With an index bid, suppliers bid how much of a discount below the index they are willing to sell the product. A *descending* bid auction is the format most commonly associated with a procurement auction - suppliers must submit increasingly lower bids to win. This bid calculation format is appropriate when the suppliers' products are fairly homogeneous and hence can be compared solely on price. A *transformational* bid is a means of standardizing bids across heterogenous suppliers. For example, the suppliers may differ depending on location, technology capabilities, and/or quality. With a transformational bid, each supplier has associated with it a transformation factor for each of the criteria deemed important in the supplier selection process. These criteria and transformational factors are agreed upon by FreeMarkets and the buyer in advance of the on-line auction, and the transformational factors are either added or multiplied to the supplier's bid, so as to create comparable bids. While each supplier is made aware of the criteria that will be taken into consideration with a transformation bid, they do not know the exact individual-supplier-level factor associated with each criterion. *NPV* bid format is typically used when the buyer and supplier will enter into a longer-term contract (e.g. 3 years), and the supplier wishes/needs to change its bid over time. For example, a buyer and supplier may enter into a contract for a good where there is learning by doing in production. NPV allows the supplier to reflect any potential cost reductions into its bid by incorporating a discount rate; the suppliers' asking price is then discounted accordingly over time.

In addition to the four auction dimensions mentioned above, FreeMarkets must decide whether or not to hold a qualifying bid round, impose individual reserve prices, the size of bid decrements, and the appropriate stopping rule. In the *qualifying round*, suppliers submit bids on individual items and/or lots via a binding one-shot sealed bid auction conducted off-line. A qualifying round is typically used when there is some uncertainty in the marketplace which FreeMarkets and the buyer were unable to resolve during the preparation period. For example, a qualifying round may be used to help FreeMarkets identify (*i*) the degree of competition for a particular product market via the level of interest, i.e., the number and level of bids submitted; (*ii*) which suppliers to invite to participate in the on-line auction; (*iii*) which items are best grouped together into a lot; and (*iii*) appropriate (possibly individual-supplier) reserve prices for the on-line auction. In selecting the size of the *bid decrements*, the auction maker must trade-off the rate at which the auction progresses versus the ability of suppliers' to reflect their true costs.[16] Generally, the auction's *stopping rule* is that only a price reduction on the lowest bid will extend the auction into overtime. However, when the lowest price is not visible to all market participants, such as in the rank feedback form, FreeMarkets may relax the stopping rule to any of the first n^{th} lowest bids can extend the auction's closing time.

3.2 Combating Collusion

One of the biggest challenge that FreeMarkets faces is designing an auction in a small and potentially collusive supplier market. In these markets, on-line and semi-transparent auctions offer a great opportunity for reducing a buyer's procurement costs. But in order for the cost savings to be realized, FreeMarkets must convince all or most eligible suppliers to participate. FreeMarkets' recent experiences with designing auctions for the ocean freight and air parcel industries illustrate the many hurdles encountered in designing auctions for potentially collusive supplier markets, with the outcome not always being a success.

FreeMarkets recently conducted an auction for a large global customer who needed ocean carrier services over routes stretching all over North America, Europe, Asia and Africa. The ocean carrier(s) would transport a variety of products for the buyer as well as handle many of the services associated with international trade, such as customs clearance. The ocean freight market is dominated by a few large (global) carriers, who handle anywhere from 30-40% of the global ocean freight business, but have much higher concentrations on a few key niche routes, and many smaller local carriers. The presence of the largest three to five Tier 1 ocean carriers was deemed as necessary for the auction's success. While the buyer would need to procure the services of multiple carriers, due to the size and scope of its demand, it wished to minimize the number of carriers with whom it had to deal.

Upon initial invitation (40 carriers were invited to participate), the majority of the Tier 1 carriers declined to participate in the on-line auction. The carriers were opposed to an on-line semi-transparent auction for multiple reasons. Their main concerns were that (*i*) the transparency of the bids on the auction would carry over and result in negative (more competitive) effects on their other business, and (*ii*) an on-line auction would result in suppliers being selected on price alone, and hence result in a commoditization of their services. In the ocean freight industry, multiple criteria factor into the performance and suitability of a supplier. Carriers may differ in the speed in which they clear customs, their turn-around time in a port, and the priority they assign a buyers' containers (when capacity is tight). The Tier 1 carriers feared that competing against Tier 2 low-cost, low-service carriers would not allow them to enjoy a justifiable price mark-up for their superior services.

By carefully designing the auction format, and speaking to the carriers individually, FreeMarkets was finally able to convince all of the Tier 1 carriers to participate. Since information transparency was one of the biggest concerns of the carriers, FreeMarkets signed a non-disclosure agreement with the carriers, stating that they would not reveal who participated in the auction. For the auction, they decided to use a rank feedback format to minimize the price information that was fed back to the marketplace and a transformation bid to take into account the service dimensions to an ocean carrier. FreeMarkets spent a considerable amount of time explaining to the carriers how a transformation bid would be calculated, and stated upfront the various criteria that would be considered in the bid calculation. All carriers were informed that their past performance with regards to the number of damaged goods, number of days of container free time, and industry-adopted service levels would be used to scale their bid. In addition, before the carriers committed to participate, FreeMarkets ran a mock (non-binding) event to demonstrate how the auction would work.

Once the carriers agreed to participate, FreeMarkets held a qualifying round to determine the supplier interest and competitive landscape on all of the lanes. The sheer number of lanes being auctioned (over 1500 lanes) complicated the auction lot decisions. While FreeMarkets had spent a couple of months reviewing the buyer's aggregate ocean traffic data over the previous year to establish traffic patterns and lanes (origin-destination points) requirements, they still did not have a solid understanding of which lanes/routes were attractive to suppliers and competitive vs. non-competitive. Previously, the buyer had delegated their freight requirement decisions to regional office managers. As a result, the buyer did not have a good understanding of its own aggregate demand requirements, and possible synergies in combining lanes. In the qualifying round, all of the lanes were made available for bidding on an individual basis. Based on the observed bids, FreeMarkets developed a lotting strategy for the various lanes and established carrier-specific reserve prices to be used in the on-line auction. The rationale behind carrier-specific prices was that using the lowest submitted bid as a general reserve price for all bidders would not take into consideration non-price factors, such as quality. As a general rule, anyone who is invited to participate in the qualifying round could participate in the on-line auction. However, if a carrier submitted a bid that was an outlier, it was not invited to participate in the on-line auction.

For the on-line auction, given the vast number of lanes being auctioned, FreeMarkets held three separate auction events; one for the North America, European and Asian lanes, respectively. Each auction event lasted two days. Within each auction, the sequencing of the lots' closing times was not believed to be of critical importance. Each lot was assigned a scheduled closing time; however, bidding on the lot could go into overtime if there was any change in the rank at the first and second level. If the rank changed, then the auction was extended by 60 seconds.

The on-line auction was deemed a success by the buyer. Although the rank feedback format and transformational bids took out much of the transparency out of the auction, they were necessary for the Tier 1 carriers to agree to participate, which, in turn, was necessary for the auction to be a success. By carefully selecting the criteria used in the transformational bids, FreeMarkets was able to bring price competition into a potentially collusive supplier base, while incorporating the numerous non-price attributes that critically impact the buyer-supplier relationship.

While removing some transparency from the auction was sufficient to persuade the top carriers in ocean freight to participate, it was not sufficient for FreeMarkets' experience with the air parcel industry. The US domestic air parcel industry is dominated by two companies, FedEx and UPS. When FreeMarkets was approached by a global buyer to design an auction for their air parcel requirements, FreeMarkets knew that the key to the auction's success would be attracting FedEx and UPS both to participate. The buyer preferred to procure services from only one company, so as to facilitate the tracking of packages. FedEx and UPS did agreed to participate in a qualifying round. To further entice both companies to participate in the on-line auction, FreeMarkets stated upfront that, if FedEx or UPS chose not to participate in the on-line auction, then their bids from the qualifying round would be used as the (general) reserve price for the on-line auction. As with the ocean freight event, FreeMarkets designed a transformational bid, rank feedback auction. Despite the buyer's large presence in the US (30% of its expenditure in the US), and FreeMarkets' reserve price policy, the buyer's business represented only a small fraction of FedEx's or UPS's total market and hence FedEx and UPS refused to participate in the on-line auction. FreeMarkets' experiences with the ocean freight and air parcel industry illustrate that, while customizing an auction to the particular event can help make it a success, an auction event can easily fail due to supplier wariness of auctions and reluctance to participate.

4. Bidder Support in Auctions - Manugistics' NetWORKS Target Pricing^{TM}

Given that the desired supplier base will participate in an auction, it is imperative that the suppliers have a solid understanding of their costs, capabilities, and the competition they face in order to make an auction an ultimate success. While FreeMarkets aids buyers in designing an auction, it is also important to help suppliers with their bid preparation process. For large and infrequent purchases, it is often worthwhile for the suppliers to hire consultants to aid them in identifying their optimal bid(s). For example, in the recent FCC spectrum auctions, many of the large telecommunication companies submitting bids hired leading game theorists and economic consulting companies to help them in their bid formation process. But for small(-er) and more frequent purchases, it behooves the suppliers to have access to a (semi-automated) bidding decision support system. As we discuss below, the need for bidding decision support is very real and is the focus of Manugistics' NetWORKS Target Pricing^{TM} software.

Manugistics, a leader in supply chain management solutions software, acquired Talus, a pricing and revenue optimization company, in 2001. The reason behind this purchase was very clear: In order to better serve their customers, Manugistics needed to be able to offer them pricing services to complement the existing planning and supply chain optimization solutions. Using advanced mathematical algorithms, Manugistics now helps customers in industrial, technology, shipping and automotive industries to determine what prices to offer to which customer segments and how to maximize profitability by managing demand in response to customer and market differences and to supply constraints. Manugistics' product, NetWORKS Target Pricing^{TM} (NTP), aids suppliers in putting together a bid for an auction. NTP was developed for use in industries where the business practice was to make price offers on a bundle of products in a contractual arrangement, one customer at a time. NTP takes into consideration the specific buyer, the buyer's order, the supplier's history with the buyer, and most importantly, competitors for the buyer's business, their previous bidding behavior and winning record when calculating an appropriate bid (or range of bids) for an auction.

NTP attempts to simplify the clients bidding problem by aiding them in systematically evaluating their *win probability* given a particular bid with a particular buyer. The decision support software assumes that if a bid is submitted, the supplier has sufficient capacity to supply the buyer's entire order and that the order utilizes only a small fraction of the supplier's capacity, i.e., NTP does not incorporate any supply constraints.[17] Rather the software focuses on determining the optimal range of prices from which the supplier should submit a bid, with each price corresponding to a different win probability.

It is a complex problem for the supplier to decide what offer to make to a particular customer. While the supplier's decision problem is theoretically simple, i.e., to determine the price offer that will maximize the expected contribution per unit = Pr(winning the bid)*(unit contribution), the difficulty lies in accurately predicting the probability of winning. Except in the case of commodity products, suppliers are almost never evaluated on price alone. Rather, a list of supplier characteristics, such as size, geographic location, whether or not the buyer has done business with the supplier before, etc., are taken into account in conjunction with price to determine which supplier is best suited for a buyer. In order to make a bidder support tool operational, a supplier must be able to estimate and maintain a parameterized win probability equation that predicts win probability as a function of the characteristics of the offer itself, the customer and the market.

This process of determining the appropriate probability functions is comprised of two distinct steps. Initially, Manugistics receives from its client a list of its (the client's) customers, previous auctions in which it participated, the bids it submitted, which of its competitors were present and their bids (if available) and the auctions' outcomes. From this list, the first step is to identify the best set of predictors of win probability and their interactions. That is, Manugistics must identify attributes of the buyer and market that are important predictors of the probability of winning and for which there exists the least intra-segment heterogeneity in outcomes (both wins and losses), and the greatest intersegment heterogeneity. These attributes can include characterization of the buyer itself, such as size, the buyer's valuation on various supplier attributes and its current supplier(s), as well as characterizations of the buyer's marketplace, such as the number of competing suppliers and their historical bidding behavior and winning record. In addition to buyer attributes, product attributes that influence the probability of winning are also identified and their effect quantified, such as the length of the contract and which package of items is being requested by the buyer.

Once the key attributes are identified, the second step is to estimate the parameters of the win probability model, using binary dependent variable regression techniques. The output of the regression model is a coefficient for each attribute, measuring the effect of that attribute on the probability of winning. Specifying the predictor attributes and estimating coefficients are undertaken at the time that the model is initialized; then periodically the steps are repeated to ensure that the model specification and parameterization remains fresh. With the repeated use of NTP, it is possible to test and refine its prediction accuracy. This is done by comparing the aggregated predictions of the model to the percentage of wins (orders actually won via an auction) actually observed. In order to make a valid comparison, the predictions are computed at the prices actually offered. By tracking differences in the predicted win probability at offered prices to the observed outcomes, the managers of the tool can detect the need to update the model coefficients.

Once NTP is installed and the prediction model is configured, a supplier can use NTP to compute the optimal bid price for a specific customer. The software has the flexibility to incorporate optional side constraints, such as constraining the price offer to be within a certain range, perhaps a range defined in terms of the price which that particular customer paid on his last contract; or constraining the win probability to be no less that some value.

NTP is currently used by a firm that supplies shipping services to businesses. The shipping firm makes tens of thousands of offers per year to thousands of different customers and prospects. There are a relatively small number of products that are offered to each customer, differing by the type of shipping (e.g. bulk versus priority documents); the destinations (e.g. domestic versus international); speed (e.g. overnight versus two-day); and ancillary services (e.g. delivery notification). The shipping firm makes over 60,000 offers per year to over 35,000 different customers and prospects. Customers differ from each other in terms of the mix of products and the volume of shipping services that they require. The also differ in terms of their own characteristics, such as the valuation that they put on reliability, their current supplier of shipping services and the competitive alternatives that are available to them. The challenge facing the shipping firm is which bid price to submit for a particular bundle of products for a particular customer.

At the time an offer is formulated, the firm has information about the customer, including the product mix and volume requested, and certain characteristics of the customer and the market. The customer characteristics include the customer's line of business, its size, and its history of buying shipping services from the firm. The market's characteristics include the firm's competitors and an estimate of their likely offers. Since the offers are private, and there is little risk of arbitrage amongst customers, the firm is essentially free to customize the price offer that it makes to each single customer, based on this information. When a salesperson makes a call on a prospect, he determines the prospect's product mix and volume requirements. These are supplied to NTP through the firm's order management system, which retrieves relevant characteristics of the prospect and the market. Based on this information, the models in NTP will return the optimal price offer to make to that customer, and ancillary information including the expected likelihood of winning the bid.

In addition to the financial benefits from offering better bids, Manugistic's customers have reported both (*i*) a dramatic reduction in the frequency of appeals to higher authority in the marketing organization by the sales staff and (*ii*) a dramatic reduction in the time it takes to quote a customer a price when using NTP. The range of prices supplied by NTP, where each price corresponding to a different win probability, provides sales agents with the flexibility and support to bid more or less aggressively for a customer's business without needing to get re-approval for price changes. Although NTP has been primarily used in sealed-bid request-for-quote settings, its input and output structure readily lend itself for use in open auction formats.

While the benefits of such a pricing tool are potentially great, there remain many challenges to its proper implementation. A critical step in determining an appropriate bid is the correct determination of the win probability function. Currently, there are large information gaps that hinder the proper specification of the probability function. Through their various implementations of NTP, Manugistics has observed time and again that companies maintain poor records of their costs and orders. In particular, poor record-keeping is prevalent when a supplier's offer loses to a competing bid. Without 'clean' data and a clear understanding of what their product is, in which product market space it lies, and who the potential competitors are, a customer-specific pricing model is severely weakened - and can potentially hurt profits.

5. Future Directions for Research in Auction Theory

Customized auctions, when designed properly, have the potential to help buyers significantly reduce their procurement costs. The growing area of research in auctions is helping us develop a better grasp on what are the critical factors when deciding on the design of the auction. There remain, however, many disconnects between the practice of designing and implementing 'real' auctions and the bulk of the auction literature.[18] These disconnects leave several interesting directions for future research: Listed below are some of the directions we believe to be the most pressing areas for research.

Product Specification A big challenge that customized auction makers face in delivering to their customer a successful sourcing experience is understanding the buyer's specifications, its effect on the potential supplier market and the optimal auction design. The bulk of the auction literature examines auctions in an environment where the product is standard and well-identified and suppliers' costs are exogenously determined. However, a buyer can influence a supplier's cost by the way that she specifies the product(s) she wishes to procure. For example, the terms of delivery, frequency of replenish, and quality of the products are all factors that can influence a supplier's cost. A serious omission from the procurement auctions literature is the important interaction between the detailing of the product, its impact on the eligible supplier base, suppliers' costs and the appropriate design of an auction.

Bundling In addition to product specification, a particularly interesting challenge posed by these B2B auctions is that they typically involve the exchange of multiple products/goods. For example, in timber, pharmaceutical, industrial parts, and electronic auctions, buyers typically wish to procure multiple units of a product, and more often than not, many different types of products via a single auction event. Unfortunately, the research in auction theory has traditionally focused on the use of single-unit auctions. One might be tempted to generalize the results of the single-object auction to the multi-object framework; however, this trust in generalizations would be misplaced. Recent work in the area of multi-object auctions has indicated that they are fundamentally different than single-object auctions and hence generalizations are invalid [Ausubel and Cramton, 2002] [Back and Zender, 1993] [Bikhchandani, 1999] [Engelbrecht-Wiggans and Kahn, 1998]. In settings where suppliers experience synergies across multiple products or units of a product, the auctioneer must decide how to bundle the units together into lots. As noted earlier, the manner in which products are bundled may affect suppliers' costs as well as their ability to bid on any particular bundle. For example, by what criterion should a buyer decide which products to bundle together; competitiveness of the supplier base, cost complementarities, and/or profitability of the product? Given the multi-unit/product nature of most B2B transactions, and potential asymmetries in suppliers' ability to compete for different products, we must gain a better understanding of the impact of product bundling on supplier bidding behavior. [19]

Transformational Factors and Scoring Rules One of the biggest challenges in designing and auction is bringing into consideration the multiple factors that influence a buyer-supplier relationship.[20] Only in special cases (e.g. purchase of a commodity product) is a buyer wise to judge a supplier on the basis of price alone. Typically, many other criteria are important in the supplier selection process; for example, a supplier's quality, reliability, size, degree of vertical integration, and available capacity. Once a buyer has identified which supplier characteristics are of utmost importance to her, she must decide on how much of her preferences to reveal to suppliers and an appropriate way to judge the suppliers on the multiple attributes. Cognizant of the multi-attribute nature of suppliers and the supplier selection process, both FreeMarkets' and Manugistics' tools try to incorporate relevant supplier attributes into the pricing decision process. There exists a small but growing literature on multidimensional auctions and scoring rules, used to transform the multiple attributes to a single number or 'score' is small [Athey and Levin, 1998] [Bichler, 2000] [Bushnell and Oren, 1994] [Chao and Wilson, 2001] [Che, 1993] [Koppius and van Heck, 2002]. Given the importance of non-price factors in the selection of suppliers, there is a great need for research that studies how the scaling and scoring of non-price attributes can affect the performance of an auction, particularly when it is difficult to observe and verify these non-price attributes.

Repeat Interaction of Suppliers and Collusion Suppliers who are part of a relatively small supplier base are bound to compete against one another over time. However, with few exceptions [Elmaghraby, 2003b] [Katzman, 1999] [Klotz and Chatterjee, 1995] [Milgrom, 2000], the auction literature almost always assumes a one-shot framework, ignoring the strategic factors that may arise in a repeated interaction framework. If an auction is to be conducted in a small supplier base, then it must take into account the presence of future auctions on a supplier's bidding behavior in the current auction and the possible presence of collusion or bidding-rings. For example, if the buyer's order constitutes a significant portion of a supplier's capacity, the supplier may not be able to participate in future auctions. Furthermore, if the suppliers are part of a bidding-ring, then they may not participate under certain auction formats (or may be uncompetitively), as experienced by FreeMarkets.[21] Therefore, an important area for future research is the design and bidding behavior of suppliers who repeatedly interact over time.

Auctions in Supply Chains An additional complicating factor in B2B auctions is that buyers and sellers typically interact across multiple channels in the supply chain. For example, a supplier may sell to a distributor (the buyer in the auction) as well as compete with the distributor in selling to the final customers. To the best of our knowledge, the research in auction theory has ignored the competing interests of different echelons in a supply chain, and its effect on the optimal design and performance of an auction.

6. Precision Pricing - Manugistics' Networks Precision Pricing

While the majority of revenue management tools continue to focus on capacity allocation, research has already begun to explore the role of simultaneous price and inventory optimization (please see [Elmaghraby and Keskinocak, 2003] [Bitran and Caldentey, 2002] for reviews of the literature). The general approach in this area of research is to assume that *any* customer arriving a particular point in time will see the same price; the question then becomes how to set price given existing inventory conditions, future demand, and the possibility (or lack thereof) of replenishment. While valuable and necessary building blocks for increasing our understanding of pricing in the presence of inventory constraints, this line of research ignores the fact that a seller may want to price discriminate across *customer types*, in addition to across time.

Many sellers operate in an environment in which they face heterogeneous customer types, i.e., their willingness-to-pay and demand elasticity are sufficiently heterogeneous, and hence wish to develop customer-type specific prices. If the seller receives buyer inquiries on a very frequent basis, an auction is a costly and time consuming pricing mechanism. In such an environment, an appropriate pricing policy should be dynamic, not incur large costs each time a transaction takes place, and should factor in the relevant characteristics of the marketplace (such as customer types) in determining the optimal prices. One pricing policy that possesses all of these characteristics is precision pricing.

Precision pricing entails simultaneously quoting each customer-type a different price based on relevant and observable characteristics and supply conditions. In order to successfully implement precision pricing, it is imperative that (*i*) the seller has some information about a customer-type's willingness to pay (*ii*) the customer market is comprised of more than one customer type (where customers' demands differ between types), (*iii*) the seller has a low-cost way to identify in which group a particular customer belongs and (*iv*) the seller is able to prevent or limit arbitrage across customer groups.

Precision pricing is a widely practiced and accepted pricing policy on the B2C front; for example, clothing stores charge a lower price at their outlet stores (charging a different price based on sales channel) and hotels offer a discounted rate to AAA members (charging a different price based on customer membership to a group). While some companies have practiced precision pricing in their B2B transactions, they have typically done so without any decision support system. Rather, they have relied on the expertise of their marketing departments to make the 'right deal' and quote the 'right price'. Given the numerous products companies typically sell, often in the thousands, the different types of customers in the marketplace, the duration of B2B contracts and the uncertainty surrounding a buyer's actual demand, a seller's pricing problem is extremely complex. Hence, the absence of decision support systems has typically left 'money on the table'. The practice of precision pricing poses several challenges, and a considerable amount of expertise is needed both in selecting a pricing policy and in periodically (or preferably continuously) evaluating its appropriateness and modifying it accordingly. Currently, Manugistics has developed a pricing solutions software, Networks Precision Pricing (P2) to aid companies in their practice of precision pricing. P2 was developed for use in industries where the business practice was to post and maintain prices that are available to any customer that qualified (as opposed to determined on a customer by customer basis).

P2, similar to traditional revenue management tools, is designed to be used in an environment where a seller experiences a stochastic arrival of heterogeneous customer orders. Both types of tools are built on the premise that there is sufficient customer heterogeneity in the buyers' market, the ability to limit arbitrage, as well as sufficient information to accurately sort customers into different segments. The main difference between P2 and revenue management tools is that P2 *explicity* incorporates the relationship between price and demand, and hence prices, not capacity allocations, are dynamically updated over time to maximize profits. Using historical information on customers' purchasing behavior, current production commitments and expected future demand, P2 helps companies develop and maintain over time prices for *each segment* of their customers types. It does so by including an optimization program (a static linear program) that determines a schedule of spot prices (one for each customer segment) that will maximize total expected profit, given each segments' estimated sensitivity to price, the expected demand and supply over the planning horizon, as well as the current (committed) production orders and the company' costs. The spot price is valid until the program is run again and an alternative optimal price schedule is determined; this updating may occur daily or weekly. Optional side constraints may be present, such as constraining the price or the expected sales to be within specified limits, or constraining the relationship between the different segment prices for a given product. In addition to supplying a price recommendation, P2 supplies price-sensitive demand forecasts and provides a means for analysts to interact with the model parameters and price recommendation.

There are two demand-side models that must be parameterized and maintained in order to operationalize the software for a particular client. The first model is the price sensitivity model which captures demand response, by market segment, to changes in each product's price. This price sensitivity model is a composite demand model that reflects general market conditions, including the presence of competitors and their influence on demand elasticity. In order to estimate price sensitivity, Manugistics and the client must first agree upon the proper customer segmentation. This is equivalent to determining how many different price points the client can support and how the client can easily 'qualify' customers for a particular price. The customer segments should reflect classes of customers whose underlying demand elasticity is different and for which eligibility can be easily determined.[22] In a supply chain setting, the size of the customers, their place in the supply chain, their end markets are all examples of observable characteristics that may correspond to different market segments. In designing the customer segments, the ability of customers to arbitrage, i.e., resell the products to customers who face a higher price, must also be taken into account.

Once the proper customer segmentation has been decided upon, the price sensitivity of each customer segment to the various products must be calculated. This is equivalent to determining the shape of each segment's demand curve for a product. In addition to price-elasticity differing across customer segments, it can also differ across products. In order to determine the demand curve, relevant market characteristics must be identified. For example, if some of the products are "near-commodities", i.e., traded in a highly competitive market, then customers will be very price sensitive. Likewise, if the product is a non-commodity, then the seller will have some market power and customers will be less price sensitive. Once the seller has selected a segmentation for their customers and products, it is crucial to validate the segmentation with historical data. Manugistics runs statistical tests to see how consistent the client's classification of customers and products is with the data. Similar to revenue management tools, P2 does not explicitly take into account the presence of competitors, but rather includes competitors presence in their estimation of demand elasticity.

The second model which must be parameterized is the time-series forecast model. This is equivalent to determining the location (placement) of the demand curve for each customer segment for each product. The placement of the demand curve is of critical importance, since it represents the potential size of the market and whether or not supply constraints will bind. If demand from a market segment is large relative to supply, then the optimal prices are adjusted to reflect the opportunity cost associated with limited supply.

One of the main challenges in correctly using P2 and implementing 3^{rd} degree price discrimination is sparsity of data. Given the typically large number of products sold by the client, many of which have very sparse transaction histories, it is not possible to obtain robust estimates of the sensitivity parameter independently for each product. In order to deal with this, P2 makes use of market structure characteristics of each product to pool transactions for the purposes of estimation. The estimates obtained from this group of products are refined for each individual product, based on reliability measures from the estimation.

In addition to properly characterizing the demand side of the market, Manugistics must help its clients correctly capture their supply side. This involves two components: (*i*) helping its clients settle on the correct definition of cost and what their *marginal* costs of production are; and (*ii*) identifying how much of each product the firm is able to sell in each period, i.e., its 'capable to promise' supply forecasts. Traditionally, suppliers have a loose understanding of what the marginal cost to produce an item is, due to inadequate cost tracking procedures. In addition, they often incorrectly factor in fixed (overhead) costs into their prices when submitting a bid. Manugistics aids their clients in properly identifying the costs that are variable, and hence the actual marginal cost associated with production. Capable to promise supply forecasts tell the seller how much of each product it can sell within each price time horizon. This quantity is important to specify when there are constrained, multi-use resources consumed in the production process. Currently, P2 takes the "capable to promise" quantity as input; this quantity is generally computed from a complementary manufacturing optimization software solution, where demand is modeled as fixed (no own or cross price sensitivity considerations).

P2 is generally used in an environment where there is a limited amount of supply; however the products are not necessarily perishable and the seller may have the opportunity to replenish supply in a timely fashion, i.e., supply is not fixed. To accurately capture the supply conditions, it is important for the client to select the proper planning horizon for pricing decisions. If the client is able to replenish its inventory quickly (e.g. within one week), then a short planning horizon whereby the client need only take into account demand in the very near future when selecting a current price, is most appropriate. However, if, as in many manufacturing settings, the time to replenish inventory is long, then the planning horizon must also be long, so as to correctly capture supply constraints and the relevant trade-offs of selling a product today versus keeping it to satisfy future demand.

P2 is currently used by a firm that supplies high-technology components to original equipment manufacturers (OEMs) and other manufacturers in the semi-conductor industry. P2 is used by the firm to maintain the price lists in its order management system. The firm sells its products through various channels, including inside and outside sales forces and a network of distributors. There are a very large number of products, on the order of 20,000 base products and upwards of 50,000 products considering special-purpose packaging. A large portion of the firm's business is derived from a few main (large) customers, with whom it negotiates prices in a contractual arrangement. The remainder of the firm's sales are derived from spot purchases.[23] Customers differ from each other in terms of the mix of the firm's product that they require for their own production and the lead time with which they place their orders. Demand is highly variable, with customers from different end-markets operating on different business cycles. Production is also global with complex supply and production processes, yielding a variable supply base as well.

The pricing problem for the firm is to decide what spot prices to post for each product and how the spot prices should be changed from one period to the next. At a point in time, the firm has information about previous demand expressed by all of its customers, certain customer characteristics, the current backlog on its order books, its own production plans, expectations regarding future demand and the characteristics of the markets in which its products are sold. Historically in the semi-conductor industry, customers were quoted prices according to their size (the total dollars spent in the semiconductor market) and whether the customer was a direct customer or a distributor. For these reasons, Manugistics's client choose to segment the market into four groups: small direct customer, medium direct customer, large direct customer and distributor. The decision to select only three size groups was a compromise between the desire to have as many groups as possible to allow for more precise pricing and the lack of sufficient data. Distributors were selected as a fourth and final group due to the different nature of their supply chain. The market characteristics used to forecast demand for each product included the *sector* in which each product is used (e.g. wireless telecommunication versus industrial process control) and the *degree of competition* that each product faces (e.g. commodity product versus exclusive, patent-protected technology). Furthermore, expectations regarding future demand were obtained from various third-party information sources and forecasts.

P2 operates in a batch mode, in which it is loaded with recent transactions and changes in supply conditions every week. Once the spot prices are decided, they are posted in the firm's order management system and become visible to the sales force and to its customers. The price sensitivity and forecast parameters are re-estimated each week with a rolling set of transactions; the results from these re-estimations are used to generate new forecasts and re-optimize the prices for every product. These new prices are reviewed by the firm's pricing analysts and, subject to their approval, loaded into the order management system. When a customer asks a salesperson for a price, the salesperson looks up the product in the order management system and, based on the segment to which the customer belongs, quotes the price. If accepted, an order is created and scheduled and its details are passed along to the fulfillment processes.

In order to do proper forecasting of demand, historical transactions are price-normalized to account for effects of price changes on observed demand. Then products are pooled based on the sectors into which they are sold and time series forecasts are generated for each sector. The sector forecasts are then distributed back to each product. A key issue that arises with the use of such a tool is quality assurance. It is important to know that the predictions that are made by the model are accurate, especially the predictions of customer response to price changes. This is assessed primarily by comparing the aggregated sales forecasts to actual sales. In order to make a valid comparison, the forecasts of sales at actual prices (as opposed to forecasts at optimal, but not yet implemented prices) are compared to observed sales. By tracking the forecast errors over time, the managers of the tool can detect the need to revise pooling rules or introduce persistent forecast adjustments.

The firm has obtained several benefits have been achieved by the use of this technology. Prior to employing P2, the firm suffered from problems of regional arbitrage. Each of its regional offices are able to quote customers pricing decisions. Given that several of their customers are large/global firms, the customers took advantage of the regional price differences. The firm used P2 as a catalyst towards centrally determined prices to help circumvent arbitrage and improve price quotes. The implementation is still new enough that there has been no comprehensive analysis of increased margin benefits, although preliminary investigation has set expectations around an increase of several percent. Additionally, the firm has been able to increase the frequency with which they evaluate and update spot prices from quarterly, at best, to weekly.

7. Future Directions for Research in Price Discrimination

Since 1920, when Pigou [Pigou, 1920] laid forth the concept of 3^{rd} degree price discrimination, (hitherto referred to as precision pricing) several questions concerning its use have been posed. Does 3^{rd} degree price discrimination increase or decrease prices? Does price discrimination always make a seller off? When should we expect to see price discrimination occur? As the practice of 3^{rd} degree price discrimination enters into various layers of a supply chain's echelons, new questions are posed. How does practicing price discrimination upstream in the supply chain differ from practicing it with only the end-customers? What effect will price discrimination have on the inventory in the supply chain? How does competition in various levels of the supply chain effect the practice of price discrimination?

While there has been a sizable amount of research done on the use of price discrimination on the B2C front, there has been relatively done on the use of price discrimination in B2B markets. One might believe that the welfare and price results for a B2C markets should carry over to B2B markets; however, the literature in this area has found this by-and-large to be not true [Katz, 1987] [DeGraba, 1990] [Yoshida, 2000].

Precision pricing, as with auctions, is posed to take advantage of the Internet's connectivity and information gathering abilities to aid companies in the B2B transactions. As companies gain more information about their customers' behavior and demand, companies can better identify and segment their customers and price accordingly. Given the apparently different nature of B2B price discrimination for B2C, it is important that we gain a better understanding of how and when it should be implemented, its benefits and its shortcomings.

Price Discrimination in a Supply Chain In B2B environments, the customers are themselves manufacturers or retailers who potentially differ along multiple attributes. In addition to factors that influence their willingness to pay, such as the end market in which they sell and the type of product they sell, these customers may also differ in (*i*) their order size, (*ii*) the specific combination or specification of products they wish to purchase and (*iii*) the urgency with which they need their order, i.e., maximum acceptable leadtime. The current literature and practice of precision pricing either assumes that the seller has a fixed quantity of product to sell (takes it as input into the model, as in the case of Manugistic's P2), or can produce instantaneously at a constant variable cost [Katz, 1987] [DeGraba, 1990] [Yoshida, 2000], an assumption seldomly satisfied in reality. A supplier's true costs are rarely comprised of only a constant variable cost. Future work on precision pricing would greatly benefit from the consideration of how pricing will affect the stream of orders a supplier receives, and in turn, how that affects the supplier's cost of meeting his demand obligations, given that a seller can (partially) observe a customer's previous purchasing history. That is, the pricing problem should be set in a multi-period horizon that incorporates possible production constraints and fixed costs. By doing so, we will be able to answer important questions such as: Should the seller develop a price schedule that will result in steady use of his capacity? Should the seller set prices so as to be able to produce his economic order quantity? How much and what type of information should the seller use about customers when selecting his prices? How far in advance should the seller accept orders (for delivery at a future date)?

In addition, further research should be done on the impact of price discrimination on the inventory in a supply chain. For example, how will downstream companies respond to different price discrimination policies? Is the seller better off implementing a static pricing policy, or should he dynamically determine the optimal prices?[24]

Price Discrimination under Competition The ability to price discriminate does not imply that a firm *should* price discriminate. Offering different prices to different customers can lead to loss of customer goodwill. One need only look at the Amazon debacle to comprehend that while tailoring prices to different market segments may theoretically boost prices, in practice it can lead to a backlash of angry customers and lost sales [25]. While this type of response to price discrimination is more likely to occur in B2C markets rather than B2B markets, it does lead to the question, 'Is price discrimination always a profitable policy for a seller?' In addition to the ill-will price discrimination can cause, it is not clear that it is always a profit boosting pricing strategy. Obviously, a monopolist who price discriminates will always increase his profits above the uniform-price profit level, for offering the same price to all customers is always an option. However, if a firm is not the only seller in the market, price discrimination may backfire and cause the firm to offer all of its customers a lower price than uniform pricing, yielding it lower profits [26]. Conversely, it may make all of the sellers better off. For example, if one were to carefully examine all of the markets in which price discrimination takes place, many of these markets could be characterized as fairly competitive, e.g. magazine and journal subscriptions, prescriptions at drugstores, restaurants, and hotels. Therefore the impact of competition on the practice of price discrimination is an important area of future research.

Behavioral Price Discrimination One of the greatest benefits of employing precision pricing on the Internet will be the ability to track and gather information about customers' purchasing behavior. Currently, there are a number of companies on the B2C front who are taking advantage of the Internet's information possibilities (e.g. Double Click and I-Behavior). The possibility for similar information tracking is just as great on the B2B front, with even larger rewards. As demonstrated in the case of Manugistic's P2 software, companies are aware that whether or not a customer has made a previous purchase will influence their willingness to pay. The reasoning behind this is typically that there are switching costs associated with moving to a new supplier or product. Once a customer has made a purchase from a particular buyer, her demand becomes more inelastic, and the seller can exploit this to increase his prices. Therefore, the optimal practice of behavior price discrimination in B2B markets is an research area worthy of immediate attention.[27]

8. Conclusion

The purpose of this chapter is to bring to the surface the current pricing practices in B2B e-marketplaces. We discussed two pricing policies in particular; *customized auctions* is currently used in e-marketplaces, *precision pricing* is a natural direction forward for many companies exploring revenue management practices in their B2B markets. Where a seller is able to use individualized prices and/or is able to effectively segment customers, we foresee large infrequent transactions remaining in the realm of auctions, while precision pricing will be employed for smaller more frequent transactions. Precision pricing and customized auctions both offers their users the promise of increase profitability. But this promise does not come without serious challenges. To better understand the design dimensions and implementation challenges of using auctions and precision pricing, we presented a few pricing 'success stories'. The challenges faced by FreeMarkets and Manugistics are not unique to these companies, but are likely to be faced by any company who wishes to implement sophisticated pricing policies in e-marketplaces. The ability of these pricing policies to improve company profitability will rest on the models used to capture and represent the market and market participant environments. We believe that the ORMS community can have a large role in improving the design and use of these and other sophisticated pricing mechanisms in e-marketplaces.

Notes

1. These prices are generally determined by a marketing department in a separate division of the company from where the revenue management tools and decisions are made.

2. This observation is based on discussions with FreeMarkets, Manugistics, and numerous other dynamic pricing software solution providers.

3. Friedman and Rust [Friedman and Rust, 1993] provide and excellent survey of the double auction literature.

4. Logistics.com was recently bought out by Manhattan Associates.

5. Another pricing mechanism appropriate in these situations is negotiation. While negotiations tend to be held between a single buyer and seller, auctions allow multiple sellers to compete for the buyer's business simultaneously, hence increasing the competitive forces during the procurement process.

6. In a combinatorial auction, bidders can submit package bids. Please see Chapters 4 and 10 for further discussion of combinatorial auctions.

7. It is important to note that a growing area of multi-unit auction literature that has been left out of the discussion below is the design and use of combinatorial auctions. These auctions, where bidders can submit package or combinational bids, are often desirable when bidders realize synergies across objects in a multi-object auction. While extremely useful in helping to capture synergies, combinatorial auctions can be quite difficult to solve for the allocation that maximizes the seller's revenue (known as the winner determination problem).

8. Presentation on March 18,2002. Data based on analysis of US Manufacturing Companies in report by CAPS, The Watch Group, A.T. Kearney.

9. Based on discussions with auction makers at FreeMarkets.

10. Krishna and Rosenthal [Krishna and Rosenthal, 1996] and Elmaghraby [Elmaghraby, 2003a] address the issue of asymmetric sized suppliers competing against each other in an auction, and its impact on the performance of an auction. Both study auctions where there are two types of bidders, global and local. Global bidders desire more than one object and their valuation for multiple objects is greater than the sum of each individual object's valuation, while local bidders desire at most one object. [Krishna and Rosenthal, 1996] identify equilibrium bidding strategies when individual objects are auctioned individually and simultaneously. [Elmaghraby, 2003a] examines under what market settings it is profitable for a buyer to exclude local suppliers from a sequential 2^{nd} price auction.

11. An alternative to the buyer explicitly bundling together products into lots and having bidders submit all-or-nothing bids for the entire lot, is to allow the bidders the flexibility to create their own bundles or packages. Such an auction format is referred to as a combinatorial or package auction, and is discussed further in Chapters 4 and 10. FreeMarkets has begun to develop such combinatorial auctions but uses them in limited market settings. Their combinatorial auction format allows bidders realtime market feedback, i.e., bidders can continuously update their bids, as opposed to discrete rounds. In addition, FreeMarkets is in the development stage with a *take-the-market* myopic best response feedback option: Based on the current bids this option tells the bidder what is the minimal bid necessary to win a particular package.

12. The issue of ordering in auction design is addressed in Benoit and Krishna [Benoit and Krishna, 1998], Elmaghraby [Elmaghraby, 2003b], Krishna [Krishna, 1993] and Pitchik [Pitchik, 1998]. Under a complete information setting, [Benoit and Krishna, 1998] examine the revenue generated by the sale of two objects that possess common values to all bidders under sequential and simultaneous auctions. [Elmaghraby, 2003b] studies the critical interplay between the sequencing of auctions and the performance of the auction in an incomplete information setting when the buyer wishes to procure two heterogenous objects. [Krishna, 1993] examines the efficiency properties of a sequential auction of capacity to an incumbent and several potential new entrants. [Pitchik, 1998] finds that, under certain equilibria (satisfying various assumptions about bidding behavior), the revenue from a sequential auction of two heterogenous objects to two bidders with private valuations only depends upon the sequence in which the goods are sold.

13. The identity of bidders is kept anonymous; rather bidders identity other bidders by some auction-given identifier, such as a number assigned to them at the start of the auction.

14. A supplier who participates in an on-line auction can be viewed as defector by a collusive supplier market and possibly subject to punishment, such retaliatory price competition.

15. In order to decrease the attractiveness of collusion, FreeMarkets has recently begun to institute a *penalty box.* If a supplier is suspected of colluding in an auction, then it is 'placed' in the penalty box and barred from participating in future auctions for some predetermined amount of time.

16. A larger bid decrement will lead to faster price changes, but at the possibly expense of excluding some suppliers' from bidding their true cost, when it falls between two bid steps.

17. In fact, the software does not incorporate any constraints, with the possible exception of restrictions imposed on the winning prices. For example, the supplier may incorporate a constraint requiring the price(s) or win probability to be above/below a certain threshold level. While supply constraints are not explicity represented in NTP, the effects of the prospective supply constraints can be incorporated as opportunity costs and factored into the optimal price determination process.

18. Klemperer [Klemperer, 1999], Krishna [Krishna, 2002], McAfee and McMillan [McAfee and McMillan, 1987] and chapters 4 and 9 in this handbook provide an excellent survey of the auction literature. In addition, Elmaghraby [Elmaghraby, 2001] provides a in-depth survey of the use of procurement auctions.

19. The issue of bundling (or lot sizing) and its impact on the performance of an auction is studied in Palfrey [Palfrey, 1983], Seshadri et al. [Seshadri et al., 1991], and Levin [Levin, 1997]. [Palfrey, 1983] explores the optimal bundling decision for a multiproduct monopolist selling via posted prices under incomplete information about buyers' valuations. [Seshadri et al., 1991] consider a model where a buyer wishes to procure a divisible object from a subset of risk-neutral suppliers via an auction. The buyer must decide into how many (equal) parts to divide her demand, thereby determining the number of suppliers from whom she wishes to procure (multi-source). Drawing from the FCC spectrum auction experience, [Levin, 1997] searches for the optimal auction design, i.e., what is the optimal manner to auction goods when there exist complementarities between two goods.

20. For further discussion of the role of multi-attribute decision making with auctions, please see Chapter 10.

21. It is worth noting that, while Manugistics' NTP does take into account the presence of competitors when formulating a bid, it does not take into consideration the repeated interactions amongst suppliers, and the possible effects one auction's outcome may have on another procurement auctions. This type of interdependence across auctions has been omitted from Manugistics' NTP, under the belief that NTP is used for relatively 'small' purchases that would not impact a supplier's ability to participate in future auctions.

22. The archetypal example of this is different prices for students, adults and senior citizens at the movie theater.

23. While P2 is used to help establish and maintain spot prices, these spot prices are sometimes used as"reference prices" in contract negotiations, particularly in contract renewal cases where current spot prices differ significantly from the spot prices in effect when the contract was initially negotiated.

24. For example, a very common form of price discrimination is coupons (or rebates). Ault et al. [Ault et al., 2000] and Gerstner and Hess [Gerstner and Hess, 1991] examine the rationale behind the use of 100% redeemable coupons, when a manufacturer sells through a retailer to his end customers. [Ault et al., 2000] find that a universal rebate can increase manufacturer profits by reducing the incentive of downstream retailers to hoard inventories when optimal wholesale price vary over time in a predictable manner. [Gerstner and Hess, 1991] assume that the end market can be segmented into high and low valuation customers. They find that 100% redeemable rebates are less costly for the manufacturer (then reducing the wholesale selling price to the retailer) when the manufacturer wishes to encourage the retailer to price so as to sell to both end market segments.

25. A very interesting website that tallies customers' responses to price discrimination by Amazon can be found at *http://www.kdnuggets.com/polls/2000/amazon_prices.htm*

26. Such a prisoner's dilemma was shown to occur in Holmes [Holmes, 1989]. [Holmes, 1989] studies a model where two firms sell differentiated products (partial substitutes) into the end market, which consists of two customer types. Both firms have the same constant marginal cost of production, sell to both markets and behave as Bertrand competitors, i.e., they compete in price, in each market. He finds examples in which the duopolists are made strictly worse off by practicing price discrimination.

27. The issue of switching cost and 'behavioral' price discrimination, whereby customers are differentiated according to their purchasing history, is a burgeoning area of research. Recent papers by Chen [Chen, 1997], Villas-Boas [Villas-Boas, 1999] [Villas-Boas, 2001] and Fudenberg and Tirole [Fudenberg and Tirole, 2000] address the issue of behavioral price discrimination, and how price discrimination across customer types effect equilibrium prices and seller profitability. Stole [Stole, 2001] provides a nice summary of this work in Chapter 6. All of these papers assume that customers differ according to (*i*) whether or not they made a purchase previously, (*ii*) whether they are new or returning customers and (*iii*) their valuations for the product. However, in practice, customers will differ along multiple attributes.

References

[AMR, 2001] AMR (2001). B2b marketplaces report, 2000-2005. *AMR Research*. August 1.

[Anthes, 1999] Anthes, G. (1999). The price had better be right. *Computerworld*, December 21. http://www.computerworld.com/cwi/story/0,1199,NAV47_STO33288,00.html.

[Athey and Levin, 1998] Athey, S. and Levin, J. (1998). Information and competition in u.s. forest service timber auctions. Working Paper, Massachusetts Institute of Technology.

[Ault et al., 2000] Ault, R., Beard, T., Laband, D., and Saba, R. (2000). Rebates, inventories, and intertemporal price discrimination. *Economic Inquiry*, 38(4):570–578.

[Ausubel and Cramton, 2002] Ausubel, L. and Cramton, P. (2002). Demand reduction and inefficiency in multi-unit auctions. Working Paper No. 96-07 University of Maryland.

[Back and Zender, 1993] Back, K. and Zender, J. (1993). Auctions of divisible goods: On the rationale for the treasury experiment. *The Review of Financial Studies*, 4:733–764.

[Barling, 2001] Barling, B. (2001). Gxs brings hosted auctions to europe. *AMR Research*. October 22.

[Benoit and Krishna, 1998] Benoit, J. and Krishna, V. (1998). Multiple-object auctions with budget constrained bidders. *Review of Economic Studies*. forthcoming.

[Bichler, 2000] Bichler, M. (2000). An experimental analysis of multi-attribute auctions. *Decision Support Systems*, 29:249–268. 3.

[Bikhchandani, 1999] Bikhchandani, S. (1999). Auctions of heterogeneous objects. *Games and Economic Behavior*, 26:193–220.

[Bitran and Caldentey, 2002] Bitran, G. and Caldentey, R. (2002). An overview of pricing models for revenue management. *Manufacturing & Service Operations Management*. to appear.

[Bushnell and Oren, 1994] Bushnell, J. and Oren, S. (1994). Bidder cost revelation in electric power auctions. *Journal of Regulatory Economics*, 6(1):5–26.

[Chao and Wilson, 2001] Chao, H. and Wilson, R. (2001). Multi-dimensional procurement auctions for power reserves: Robust incentive-compatible scoring and settlement rules. Working Paper, Stanford University.

[Che, 1993] Che, Y. (1993). Design competition through multidimensional auctions. *RAND Journal of Economics*, 24:668–681.

[Chen, 1997] Chen, Y. (1997). Paying customers to switch. *Journal of Economics and Management Strategy*, 6:887–897.

[DeGraba, 1990] DeGraba, P. (1990). Input market price discrimination and the choice of technology. *American Economic Review*, 80(5):1246–1253.

[Economist, 2001] Economist (2001). B2b exchanges : The container case. *The Economist*. October 21.

[Elmaghraby, 2001] Elmaghraby, W. J. (2001). Supply contract competition and sourcing policies. *Manufacturing & Service Operations Management*, 2(4):350–371.

[Elmaghraby, 2003a] Elmaghraby, W. J. (2003a). The impact of asymmetric bidder size on auction performance: Is more bidders always better? Working Paper, Georgia Institute of Technology.

[Elmaghraby, 2003b] Elmaghraby, W. J. (2003b). The importance of ordering in sequential auctions. *Management Science*, 49(5):673–682.

[Elmaghraby and Keskinocak, 2003] Elmaghraby, W. J. and Keskinocak, P. (2003). Dynamic pricing in the presence of inventory considerations: Research overview, current practices and future directions. *Management Science*. Special Issue on E-Business, forthcoming.

[Engelbrecht-Wiggans and Kahn, 1998] Engelbrecht-Wiggans, R. and Kahn, C. (1998). Multi-unit auctions with uniform prices. *Economic Theory*, 12:227–258.

[Friedman and Rust, 1993] Friedman, D. and Rust, J. (1993). *The Double Auction Market: Institutions, Theories and Evidence*. Addison Wesley Longman. SPI Studies in the Sciences of Complexity.

[Fudenberg and Tirole, 2000] Fudenberg, D. and Tirole, J. (2000). Customer poaching and brand switching. *RAND Journal of Economics*, 31(4):634–657.

[Gerstner and Hess, 1991] Gerstner, E. and Hess, J. (1991). A theory of channel price promotions. *American Economic Review*, 81(4):872–886.

[Holmes, 1989] Holmes, T. (1989). The effects of third-degree price discrimination in oligopoly. *American Economic Review*, 79(1):244–250.

[Katz, 1987] Katz, M. (1987). The welfare effects of third-degree price discrimination in intermediate good markets. *American Economic Review*, 77(1):154–167.

[Katzman, 1999] Katzman, B. (1999). A two stage sequential auction with multi-unit demands. *Journal of Economic Theory*, 86:77–99.

[Klemperer, 1999] Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13(3):227–286.

[Klotz and Chatterjee, 1995] Klotz, D. and Chatterjee, K. (1995). Dual sourcing in repeated procurement competitions. *Management Science*, 41(8):1317–1327.

[Koppius and van Heck, 2002] Koppius, O. and van Heck, E. (2002). Information architecture and electronic market performance in multidimensional auctions. Working Paper Erasmus University.

[Krishna, 1993] Krishna, K. (1993). Auctions with endogenous valuations: The persistence of monopoly revisited. *American Economic Review*, 83(1):147–160.

[Krishna, 2002] Krishna, V. (2002). *Auction Theory*. Academic Press, San Diego.

[Krishna and Rosenthal, 1996] Krishna, V. and Rosenthal, R. (1996). Simultaneous auctions with synergies. *Games and Economic Behavior*, 17:1–31.

[Latham, 2000] Latham, S. (2000). Independent trading exchanges–the next wave of b2b e-commerce. *AMR Research*. March 1.

[Levin, 1997] Levin, J. (1997). An optimal auction for complements. *Games and Economic Behavior*, 18:176–192.

[McAfee and McMillan, 1987] McAfee, P. and McMillan, J. (1987). Auctions and bidding. *Journal of Economic Literature*, 15:699–738.

[McGill and van Ryzin, 1999] McGill, J.I. and van Ryzin, G.J. (1999). Revenue management: Research overview and prospects. *Transportation Science*, 33:233–256.

[Milgrom, 2000] Milgrom, P. (2000). *A Theory of Auctions and Competitive Bidding II*. Edward Elgar Publishing, Inc. in The Economic Theory of Auctions P. Klemperer Ed. Cheltenham, UK.

[Palfrey, 1983] Palfrey, T. (1983). Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica*, 51(2):463–483.

[Phillips, 2003] Phillips, R. (2003). *Pricing and Revenue Optimization*. Stanford University Press, forthcoming, Palo Alto, CA. Chap. 4.

[Pigou, 1920] Pigou, A. (1920). *The Economics of Welfare*. Macmillan, London.

[Pitchik, 1998] Pitchik, C. (1998). Budget-constrained sequential auctions with incomplete information. Working Paper University of Toronto.

[Prouty, 2001] Prouty, K. (2001). Automotive and heavy equipment outlook: Gm tries to move covisint beyond auctions. *AMR Research*. October 31.

[Rosenblum, 2002] Rosenblum, P. (2002). Major tidal wave of auctions coming to apparel, soft home, and footwear retail. *AMR Research*. August 5.

[Seshadri et al., 1991] Seshadri, S., Chatterjee, K., and Lilien, G. (1991). Multiple source procurement competition. *Management Science*, 10(3):246–263.

[Stole, 2001] Stole, L. (2001). Price discrimination in competitive environments. Working Paper, Graduate School of Business, University of Chicago.

[Villas-Boas, 1999] Villas-Boas, J. (1999). Dynamic competition with customer recognition. *RAND Journal of Economics*, 30(4):604–631.

[Villas-Boas, 2001] Villas-Boas, J. M. (2001). Price cycles in markets with customer recognition. Working Paper, Haas School of Business, University of California at Berkeley.

[Yoshida, 2000] Yoshida, Y. (2000). Third-degree price discrimination in input markets: Output and welfare. *American Economic Review*, 90(1):240–246.